

CLAIMS

1. A method of generating executable code for a configurable microprocessor architecture whereby:

there are individual registers in the data paths between execution units;

the update of these registers is explicitly managed from the instruction set of the processor; and

the reading and writing of those registers is represented in a data flow graph representation of the program being targeted to the processor.
2. The method according to claim 1 whereby the number of registers associated with particular execution units is configurable.
3. The method according to claim 1 whereby the target architecture is specified in an input file.
4. The method according to claim 1 whereby certain units may only be connected to a subset of other execution units in the architecture.
5. The method according to claim 4 whereby there is a central register file but only a subset of the execution units in the system can directly communicate with it.
6. The method according to claim 3 whereby the input program is initially represented as a sequence of operations that can be performed by execution units present in the target architecture.
7. The method according to claim 6 whereby the inputs to and results from these basic instructions may be communicated via a central register file.
8. The method according to claim 7 whereby the code sequence may be optimised to reduce the communication required with the central register file.

9. The method according to claim 8 whereby multiple reads of a given register file value may be transformed into a single read with suitable communication of the same data to other consumers of the data value.
10. The method according to claim 8 whereby pairs of central register file writes and reads may be transformed to use direct communication between the generating and consuming operations.
11. The method according to claim 1 whereby the data flow graph is transformed so that data edges correspond to physical connections in the architecture.
12. The method according to claim 11 whereby additional nodes may be inserted into the graph to represent the copying of data values where there is no physical connection corresponding to the graph data flow.
13. The method according to claim 12 whereby the dependencies between reads and writes to registers are represented as edges in the graph.
14. The method according to claim 13 whereby graphs that cannot be scheduling are detected by the presence of cycles in the graph.
15. The method according to claim 1 whereby an idealised form of the graph is generated that assumes the availability of unrestricted connectivity in the architecture.
16. The method according to claim 15 whereby the idealised form of the graph is used to influence the binding of operations to physical execution units in the architecture.
17. The method according to claim 1 whereby special edges within the graph represent communication of data via a central register file.
18. The method according to claim 17 whereby operations from different basic blocks may be represented in a single graph.
19. The method according to claim 1 whereby individual operations in the graph are bound to particular execution unit instances.

20. The method according to claim 19 whereby the unit binding uses an estimate of the delay caused by transporting operands to and results from the operation as a factor in the allocation.
21. The method according to claim 20 whereby the transport cost is dependent on the structure of connectivity between the operations in the graph.
22. The method according to claim 1 whereby the graph may be updated as new physical paths are added to the architecture in order to reduce the graph height to allow shorter code schedules.
23. The method according to claim 1 whereby individual execution units are controlled by particular bits within the overall execution word.
24. The method according to claim 23 whereby certain bits in the execution word may be used to control more than a single execution unit.
25. The method according to claim 24 whereby the allocation of bits in the execution word to particular execution units is optimised as the architecture is generated.
26. The method according to claim 25 whereby the usage of individual execution units is used to influence the allocation of the execution word.
27. A microprocessor configured to execute code that has been generated using the method of any preceding claim 1 – 26.